Forecasting transitions in systems with high dimensional stochastic complex dynamics: A Linear Stability Analysis of the Tangled Nature Model

Andrea Cairoli*
School of Mathematical Sciences, Queen Mary, University of London, Mile End Road, E1 4NS, UK

Duccio Piovani† and Henrik Jeldtoft Jensen‡
Centre for Complexity Science and Department of Mathematics, Imperial College London, South Kensington Campus, SW7 2AZ, UK

arXiv:1407.5024v2 [nlin.AO] 24 Jul 2014

We propose a new procedure to monitor and forecast the onset of transitions in high dimensional complex systems. We describe our procedure by an application to the Tangled Nature model of evolutionary ecology. The quasi-stable configurations of the full stochastic dynamics are taken as input for a stability analysis by means of the deterministic mean field equations. Numerical analysis of the high dimensional stability matrix allows us to identify unstable directions associated with eigenvalues with positive real part. The overlap of the instantaneous configuration vector of the full stochastic system with the eigenvectors of the unstable directions of the deterministic mean field approximation is found to be a good early-warning of the transitions occurring intermittently.

Introduction - Many complex high dimensional systems are characterised by intermittent dynamics, where relatively long quiescent periods are interrupted by sudden and quick bursts of activity during which the system suffers hectic rearrangements. These rearrangements can be seen as transitions between metastable states. Examples of abrupt transitions have been identified in a broad range of systems [1]: in biological ecosystems [2, 3] transitions from a flourishing to a wild state can occur, in financial markets [4] endogenous crisis can destabilize an existing balance, in the human brain [5] epileptic seizures signals a switch from a regular to an irregular condition, climate [6] can exhibit sudden changes both overall or in one of its subsystems, like when a bloom of harmful algae suddenly forms in the sea [7]. Due to their widespread occurrence, these transitions have gathered a huge interest in the last decade, with research mainly focused on developing statistical methods to forecast them from the observed time series [8–10] and on the development of a general mathematical framework to describe them [11]. In the present paper we contribute to both efforts by developing a mathematical analysis by use of a paradigmatic model exhibiting intermittent stochastic evolution and by identifying systemic observables that can deliver early-warning of impending transitions. We focus on the Tangled Nature (TaNa) model [12–14] of evolutionary ecology. The initial aim of the model was to establish a sound and simple mathematically framework for "punctuated equilibrium", i.e. the observed intermittent mode of macro-evolution.

The TaNa model is an individual based stochastic model of coevolution. The model's phenomenology is in good agreement with biological observations [15]. At the microscopic level of individuals the dynamics is unfolding at a smooth constant pace: agents reproduce, mutate and die at essentially constant rates. On the contrary, at the systemic level the generated ecological network structures jump from one metastable configuration to another (denoted quasi-Evolutionary Stable Strategies or qESS). We investigate these macroscopic instabilities by performing a Linear Stability Analysis (LSA) of the mean field representation of the dynamics about the actual configurations produced by the full stochastic dynamics. LSA is obviously a standard procedure to analyze the nature of fixed points for deterministic autonomous equations of motion. Here we develop the method to allow applications to high dimensional *stochastic* dynamics.

We recall that the LSA for a deterministic autonomous system of equations: $\dot{\mathbf{n}} = \mathbb{F}(\mathbf{n})$ for the time dependent vector $\mathbf{n}(t)$ consists in first identifying the fixed points $\mathbf{n}^*$, i.e. the time independent solutions: $\dot{\mathbf{n}}^* = 0$ of the full non-linear set of equations of motion. One next studies to the first order terms the time dependence of the deviation: $\delta\mathbf{n}(t) = \mathbf{n}(t) - \mathbf{n}^*$ about each of these fixed points: $\delta\dot{\mathbf{n}} = \mathbb{M}[\mathbf{n}^*]\delta\mathbf{n}$. This equation has the solution $\delta\mathbf{n}(t) = e^{(t-t_0)\mathbb{M}[\mathbf{n}^*]}\,\delta\mathbf{n}_0$ and the stability of a given fixed point $\mathbf{n}^*$ is now determined by the properties of the spectrum of the matrix $\mathbb{M}[\mathbf{n}^*]$, with unstable directions being connected to eigenvalues with positive real parts.

When we apply this procedure to very high dimensional situations like the TaNa model it is not possible to solve directly the fixed point equation: $\mathbb{F}(\mathbf{n}^*) = 0$. Instead, we can use the observed qESS configurations generated by the full stochastic dynamics to approximate $\mathbf{n}^*$ and perform a LSA of the mean field dynamics about these configurations. To our knowledge this procedure for applying LSA to high dimensional stochastic dynamics has hardly been attempted before. Only recently, LSA of agent-based models have been studied: in [16] the stability properties of the attractors of a generalized Sznajd model are derived from its mean-field formulation, whereas in [17] a similar analysis has been done for a net-

work of pulse-coupled neurons. Neither of these systems, however, exhibits intermittent behavior as we observe in the TaNa.The intermittent dynamics allows us to define a new mean-field based early-warning measure for the occurrence of abrupt transitions.

The model - In the TaNa, an agent is represented by a sequence of binary variables with fixed length L [18], denoted as $\mathbf{S}^a = (S_1^a, \ldots, S_L^a)$, where $S_i^a = \pm 1$. Thus, there are 2^L different sequences, each one represented by a vector in the genotype space: $\mathcal{S} = \{-1, 1\}^L$. In a simplistic picture, each of these sequences represents a genome uniquely determining the phenotype of all individuals of this type. We denote by $n(\mathbf{S}^a, t)$ the number of individuals of type $\mathbf{S}^a$ at time t and the total population is $N(t) = \sum_{a=1}^{2^L} n(\mathbf{S}^a, t)$. We define the distance between different genomes $\mathbf{S}^a$ and $\mathbf{S}^b$ as the Hamming distance: $d_{ab} = \frac{1}{2L}\sum_{i=1}^{L} |S_i^a - S_i^b|$. A time step is defined as a succession of one annihilation and of one reproduction attempt. During the killing attempt, an individual is chosen randomly from the population and killed with a probability p_{kill} constant in time and independent on the type. During the reproduction process, a different randomly chosen individual $\mathbf{S}^a$ successfully reproduces with probability: $p_{off}(\mathbf{S}^a, t) = \frac{\exp(H(\mathbf{S}^a,t))}{1+\exp(H(\mathbf{S}^a,t))}$, which depends on the occupancy distribution of all the types at time t via the weight function:

$$H(\mathbf{S}^a, t) = \frac{k}{N(t)} \sum_{\mathbf{S}^b \in \mathcal{S}} \mathbf{J}(\mathbf{S}^a, \mathbf{S}^b) n(\mathbf{S}^b, t) - \mu N(t). \quad (1)$$

In Eq. (1), the first term couples the agent $\mathbf{S}^a$ to one of type $\mathbf{S}^b$ by introducing the interaction strength $\mathbf{J}(\mathbf{S}^a, \mathbf{S}^b)$, whose values are randomly distributed in the interval $[-1, +1]$. For simplification and to emphasize interactions we here assume: $\mathbf{J}(\mathbf{S}^a, \mathbf{S}^a) = 0$. The parameter k scales the interactions strength and μ can be thought of as the carrying capacity of the environment. An increase (decrease) in μ corresponds to harsher (more favourable) external conditions. The reproduction is asexual: the reproducing agent is removed from the population and substituted by two copies $\mathbf{S}_1^a$ and $\mathbf{S}_2^a$, which are subject to mutations. A single mutation changes the sign of one of the genes: $S_i^\gamma \to -S_i^\gamma$ with probability p_{mut}. Similarly to a Monte Carlo sweep in statistical mechanics, the unit of time of our simulations is a *generation* consisting of $N(t)/p_{kill}$ time steps, i.e. the average time needed to kill all the individuals at time t. These microscopic rules generate intermittent macro dynamics. The system is persistently switching between two different modes: the qESS states and the transitions separating them. The qESS states are characterized by small amplitude fluctuations of N(t) and stable patterns of occupancies of the types (Fig. 1, respectively top and bottom panel). However, these states are not perfectly stable and configurational fluctuations may trigger an abrupt transition to a different qESS state. The transitions consist of adaptive random walks in the configuration space while searching for a new metastable configuration and are related to high amplitude fluctuations of N(t).

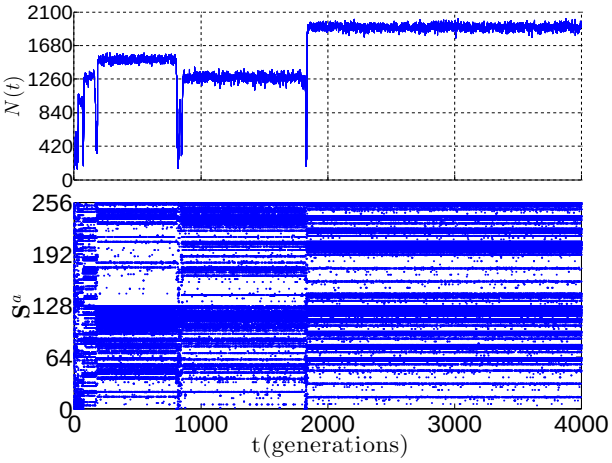

Figure 1. Top panel: total population as a function of time (in generations) for a single realization of the TaNa. The punctuated dynamics is clearly visible: quasi-stable periods alternate with periods of hectic transitions, during which $N(t)$ exhibits large amplitude fluctuations. Bottom panel: occupancy distribution of the types. The genotypes are labelled arbitrarily and a dot indicates a type which is occupied at the time t.

Linear stability analysis - We now describe the mean field deterministic approximation of the TaNa. The macroscopic configuration is given by $\mathbf{n}(t) = (n_1(t), \ldots, n_{2^L}(t)) \in (\mathbb{N} \cup \{0\})^{2^L}$ and evolves according to the continuous time mean field equation

$$\frac{\mathrm{d}\mathbf{n}(t)}{\mathrm{d}t} = \frac{1}{N(t)} \mathbb{T}[\mathbf{n}(t)]\mathbf{n}(t), \quad (2)$$

where

$$\{\mathbb{T}[\mathbf{n}(t)]\}_{ab} = (p_{off}^b(\mathbf{n}(t))(2p_0 - 1) - p_{kill})\,\delta_{ab} + 2\,p_{off}^b(\mathbf{n}(t))(p_{mut})^{Ld_{ab}}(1 - p_{mut})^{L(1-d_{ab})}(1 - \delta_{ab}). \quad (3)$$

Although being an approximation of the real stochastic dynamics [19], we find that Eqs. (2) is able to capture the low frequency variation of the discrete time model.

As mentioned above we cannot directly derive the non-trivial solutions of the fixed points' equation: $\mathbb{T}[\mathbf{n}^*]\mathbf{n}^* = 0$, so instead we use the approximated configuration $\bar{\mathbf{n}}_{stoc} \simeq \mathbf{n}^*$, which is obtained from the full stochastic dynamics as a time average during a qESS. Due to the fluctuations occurring during a qESS (See Fig. 1) $\bar{\mathbf{n}}_{stoc}$ is only approximately stationary: $T[\bar{\mathbf{n}}_{stoc}]\bar{\mathbf{n}}_{stoc} \simeq 0$. The linearized equation for the deviations away from $\bar{\mathbf{n}}_{stoc}$ becomes: $\frac{\partial}{\partial t}\delta\mathbf{n} = \mathbb{M}[\bar{\mathbf{n}}_{stoc}]\,\delta\mathbf{n}$ and the stability of a qESS is given by the spectrum of eigenvalues of $\mathbb{M}[\bar{\mathbf{n}}_{stoc}]$. How effective an unstable eigendirection is in destabilising the configuration $\mathbf{n}(t)$ will depend on the overlap between the

deviation from $\bar{\mathbf{n}}_{stoc}$: $\delta\mathbf{n}(t) = \mathbf{n}(t) - \bar{\mathbf{n}}_{stoc}$ and the unstable directions. During the qESS we therefore introduce the following instability indicator:

$$Q(t) = \max_{\lambda \in Sp^{+}(\mathbb{M}[\bar{\mathbf{n}}_{stoc}])} \left| e^{\lambda} \langle (\mathbf{n}(t) - \bar{\mathbf{n}}_{stoc}), \mathbf{e}_{\lambda} \rangle \right| \qquad (4)$$

where the eigenvalues λ and the correspondent eigenvectors $\mathbf{e}_{\lambda}$ of $\mathbb{M}[\bar{\mathbf{n}}_{stoc}]$ can be computed numerically for high dimensions (in our case with the Intel DGEEV routine). $Sp^{+}(\mathbb{M}[\bar{\mathbf{n}}_{stoc}])$ refers to the eigenvalues with positive real part and the brackets denotes the scalar product. For the TaNa we verified numerically that the stable and the unstable sub-spaces are orthogonal. $Q(t)$ simply measures the maximal expected growth of $\delta\mathbf{n}(t)$ during the time interval $\Delta t = 1$.

Procedure and Results - We monitor the system in real time. To understand when a qESS is established, we average the occupation vector $\mathbf{n}(t)$ over time windows of $\Delta T = 100$ time units (i.e. generations in the case of the TaNa) to obtain $\bar{\mathbf{n}}_{stoc}$. We then check if the system is stationary, i.e. $T[\bar{\mathbf{n}}_{stoc}]\bar{\mathbf{n}}_{stoc} \simeq 0$, repeating the process until the condition is satisfied. Once a qESS has been reached we linearize about the configuration $\bar{\mathbf{n}}_{stoc}$ and compute $Q(t)$ and the instantaneous deviation: $\|\delta\mathbf{n}(t)\| = \|\mathbf{n}(t) - \bar{\mathbf{n}}_{stoc}\|$. Below we demonstrate that the indicator $Q(t)$ is able to monitor and even forecast the onset of a transition out of the current qESS. The transition shows up directly as an unbounded sudden growth of $\|\delta\mathbf{n}(t)\|$. Once a transition out of the current qESS has occurred, we average again $\mathbf{n}(t)$ to establish the new quasi stable configuration $\bar{\mathbf{n}}_{stoc}$.

In Fig.2 we show Q as a function of the microscopic time steps (blue curve). We observe that $\|\delta\mathbf{n}(t)\|$ fluctuates during the qESS. In contrast Q only grows when a transition is about to occur. Typically Q starts to increase several generations prior to the transition corresponding, in this particular case, to thousands of single update events.

To understand the relation between $\|\delta\mathbf{n}\|$ and Q we show in Fig. 3 the joint probability density $P(\|\delta\mathbf{n}(t^* - \tau)\|, Q(t^* - \tau))$ for τ generations before the time t^* of the transition. We identify t^* as the time when the condition: $\|\delta\mathbf{n}(t^*)\| > d$ holds persistently for at least the next 10 consecutive generations for a fixed threshold $d = 150$ corresponding to the typical upper bound for the amplitude of the fluctuations inside the qESSs. From the way the region of largest support move in the $Q - \|\delta\mathbf{n}\|$ plane as the transition is approached we see to what extent monitoring Q allows one to predict the transition. Note that a significant support for values of Q larger than about 10 starts to develop from around $\tau = 5$. At these times the deviation $\|\delta\mathbf{n}\|$ is still most often below the inherent qESS fluctuation level of 150. We may encounter situations where Q gives a false signal, by increasing significantly in correspondence to small amplitude perturbations of $\mathbf{n}(t)$. Remarkably, we can see that these events happen

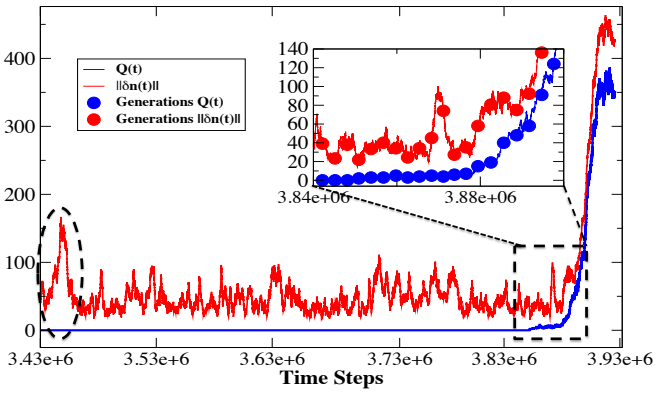

Figure 2. Typical behavior of $Q(t)$ and $\|\delta\mathbf{n}(t)\|$ in a single run of the TaNa in time steps. Clearly $Q(t) \simeq 0$ even for more rare strong fluctuations (dashed circle) inside the qESSs, whereas it starts to increase rapidly before the actual transition. In the inset, we zoom on the transition and indicate with markers the points observed at the coarse-grained level of generations. Notice that between two generations many time steps (events) are present.

with low probability, thus not affecting significantly the performance of the Q measure.

Finally, our success rate in predicting transitions is approximately 85-87%, however non-predicted transitions do occur and are related to a non-vanishing probability that a direction which is weakly stable (negative eigenvalues close to zero) of the mean field can trigger a transition. This is shown in Fig. 4, where the distribution of the real parts of the eigenvalues responsible for the transitions is plotted. This purely stochastic phenomenon explains why we find with non vanishing probability transitions together with $Q \simeq 0$ (see Fig. 3, panel (a)).

Conclusions - We combined deterministic mean field analysis with stochastically generated configurations to develop a measure capable of forecasting abrupt transitions in the TaNa model. We believe that the procedure outlined here can be applied to other high dimensional complex systems, including for example economy or neuronal systems, when sufficient data sampling is possible to establish the effective interaction matrix (the $\mathbf{J}$ matrix above) of the fluctuating dynamics. Its mean field approximation can then allow the monitoring of Q. With the rapid development towards big-data sampling capacity in many areas of science this scenario becomes increasingly a possibility.

Acknowledgments - This work was supported by the European project CONGAS (Grant FP7-ICT-2011-8-317672)

Contributions - A.C. and D.P. contributed equally to this work.

* a.cairoli@qmul.ac.uk
† duccio.piovani@gmail.com
‡ h.jensen@imperial.ac.uk

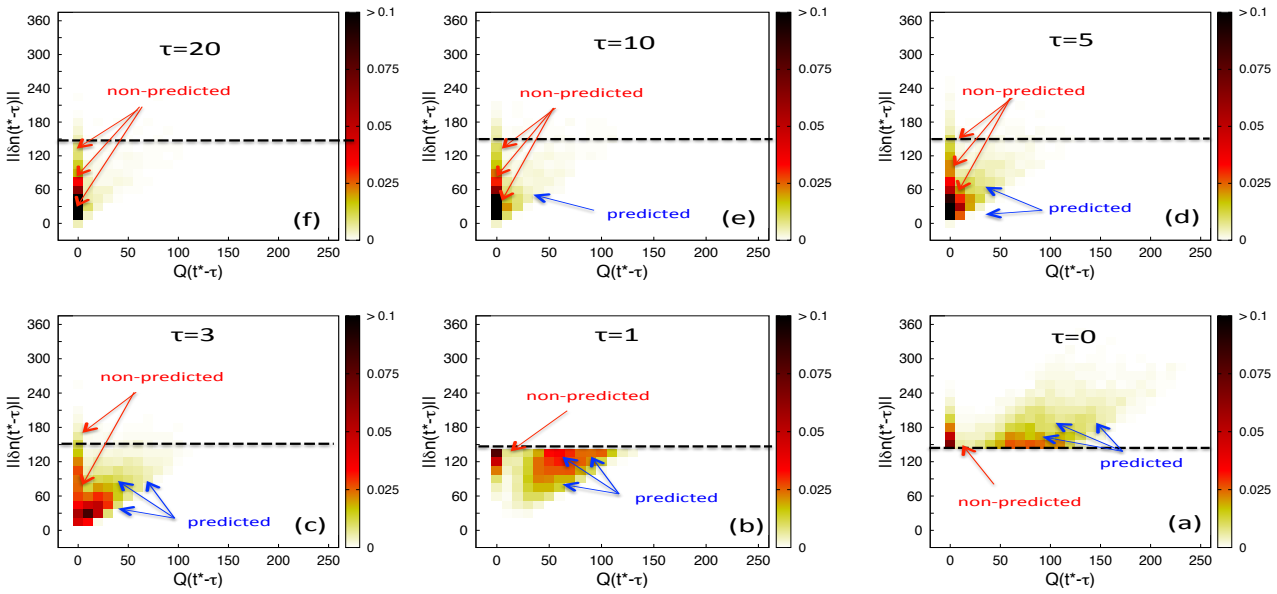

Figure 3. 2D distribution $P(\|\delta\mathbf{n}(t^*-\tau)\|, Q(t^*-\tau))$ averaged over 13000 transitions for different values of τ. The predictive power of Q is evident: typical fluctuations inside the qESSs are not signaled by Q (panels (e-f)), whereas dangerous perturbations leading to a transition are recognized by the increasing of Q away from zero (panels (a-d)). This is already seen for $\tau = 5$, which is still remarkably far from the transition. Examples of predicted/non predicted transitions are then shown with arrows in panels (d-a). The other plots can be interpreted in a similar way.

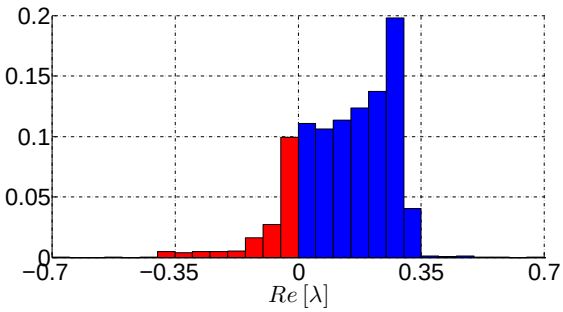

Figure 4. Distribution of the real part (red/blue boxes for negative/positive one) of the eigenvalues correspondent to eigendirections with maximum overlap with $\delta\mathbf{n}(t)$ at the beginning of a transition. The distribution is clearly dominated by the unstable eigenspace, but a significant probability ($\approx 17\%$) of weak stable eigenvalues is found.

[1] M. Scheffer, *Critical transitions in nature and society* (Princeton University Press, 2009).

[2] M. Scheffer, S. Carpenter, J. A. Foley, C. Folke, and B. Walker, Nature **413**, 591 (2001).

[3] M. Scheffer and S. R. Carpenter, Trends in Ecology & Evolution **18**, 648 (2003).

[4] R. M. May, S. A. Levin, and G. Sugihara, Nature **451**, 893 (2008).

[5] B. Litt and J. Echauz, The Lancet Neurology **1**, 22 (2002).

[6] T. M. Lenton, H. Held, E. Kriegler, J. W. Hall, W. Lucht, S. Rahmstorf, and H. J. Schellnhuber, Proceedings of the National Academy of Sciences **105**, 1786 (2008).

[7] K. G. Sellner, G. J. Doucette, and G. J. Kirkpatrick, Journal of Industrial Microbiology and Biotechnology **30**, 383 (2003).

[8] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. Van Nes, M. Rietkerk, and G. Sugihara, Nature **461**, 53 (2009).

[9] S. J. Lade and T. Gross, PLoS computational biology **8**, e1002360 (2012).

[10] M. Scheffer, S. R. Carpenter, T. M. Lenton, J. Bascompte, W. Brock, V. Dakos, J. van de Koppel, I. A. van de Leemput, S. A. Levin, E. H. van Nes, et al., Science **338**, 344 (2012).

[11] C. Kuehn, Physica D: Nonlinear Phenomena **240**, 1020 (2011).

[12] K. Christensen, S. A. di Collobiano, M. Hall, and H. J. Jensen, J. Theor. Biol. **216**, 73 (2002).

[13] M. Hall, K. Christensen, S. A. di Collobiano, and H. J. Jensen, Phys. Rev. E **66** (2002).

[14] E. P. Anderson and H. J. Jensen, J. Theor. Biol. **232**, 551 (2005).

[15] S. Laird, D. Lawson, and H. J. Jensen, in *In Mathematical Modeling of Biological Systems. Vol. 2*, edited by A. Duetsch, L. Brusch, H. Bryne, G. de Vries, and H. Herzel (Birkauser, 2007), pp. 49–62.

[16] A. M. Timpanaro and C. P. Prado, Physical Review E **86**, 046109 (2012).

[17] S. Olmi, A. Politi, and A. Torcini, The Journal of Mathematical Neuroscience (JMN) **2**, 1 (2012).

[18] P. Higgs and B. Derrida, J. Mol. Evolution **35**, 454 (1992).

[19] R. Durrett and S. Levin, Theoretical population biology **46**, 363 (1994).